USC Viterbi
School of Engineering

ZettaFlop/s

Workshop on the Frontiers of Extreme Computing 2005

Information Sciences Institute, Computational Sciences Division

8 March 2006

Bob Lucas
rflucas@isi.edu

USC Viterbi School of Engineering

Acknowledgements

ISI Information Sciences Institute

Organized by Sandia's Erik DeBenedictis

- **Ably assisted by Prof. Sterling**

Third in series:

- **Frontiers of Extreme Computing, Oct. 2005**
- **The Path to Extreme Supercomputing, Oct. 2004**

Santa Cruz, CA, Oct. 2005

- **Chaminade is a wonderful conference site**

USC

USC Viterbi School of Engineering

Topics Discussed

ISI Information Sciences Institute

- Applications
- Technology
  - Future of Moore's Law
  - Alternatives
- Architecture
- Software

USC

USC Viterbi School of Engineering

Applications

ISI Information Sciences Institute

There need to be some!

- I believe there are ☺

Applications presented:

- Climate (Bader, LLNL)
- Fusion (Jardin, Princeton PPL)
- Space (Biswas, NASA ARC)
- Graphs (Hendrickson, Sandia)

Reports presented:

- ScaLES (Keyes, Columbia)
- NRC Future of Supercomputing (Dally, Stanford)

USC

Si

ITRS Roadmap (Zeitzoff, Sematech)

Future Technologies (Theis, IBM)

Bottom line:

- **Roadmap extends to 2020 and 10nm devices**
- **Clock rates will grow ~15% per year**
- **The bar is set very high for any alternative!**

USC Viterbi School of Engineering

6nm FET

ISI Information Sciences Institute

IBM Research

IBM

… and in the lab.

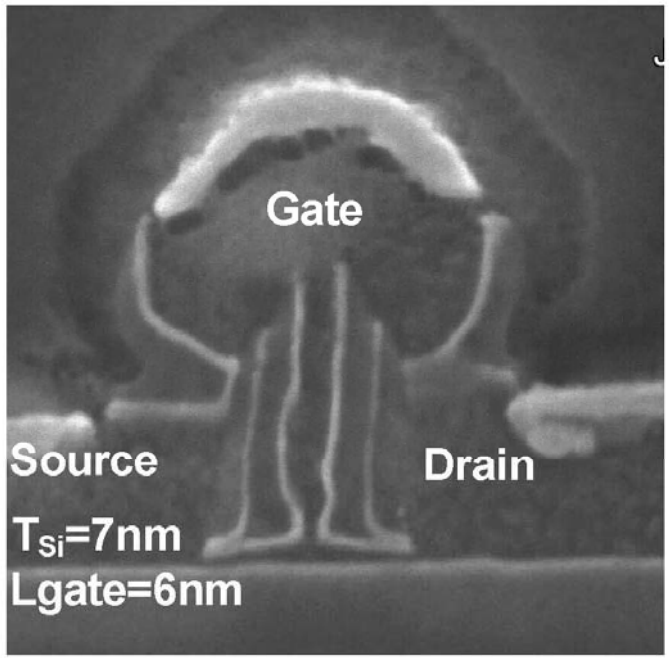

B. Doris et al., *IEDM* , 2002

USC

Slide from IBM's Thomas Theis

USC Viterbi School of Engineering

Alternative Technology

ISI Information Sciences Institute

- Single Flux Quantum (Silver, TRW retired)
- Reversible Computing (Frank, Florida)
- Nanoscale (Williams, HP)
- Nanowires (deHon, Caltech)
- Quantum Dots (Lent, ND + Neimer, GaTech)
- Quantum Computing (Foster, NSF + Oskin, Washington + Williams, JPL)

USC

USC Viterbi School of Engineering

World's densest memory

ISI Information Sciences Institute

Atomic Force Microscope Image of Prototype HP Device

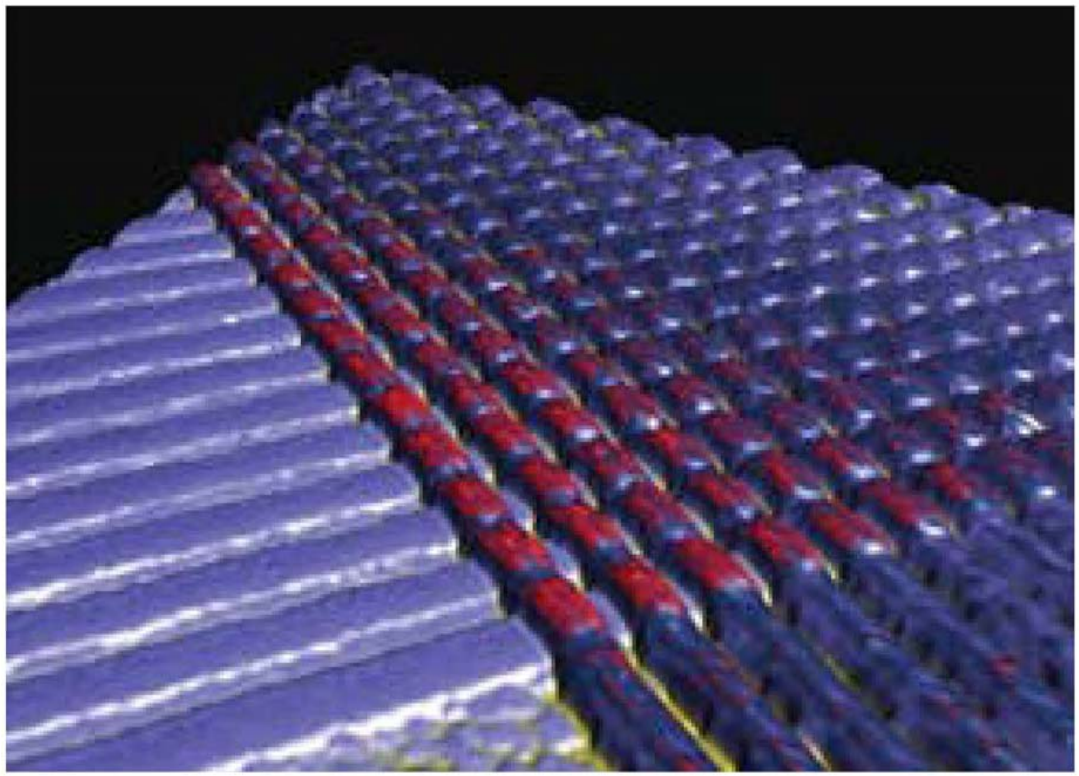

Image from DeBenedictis, who got it from Williams

USC

USC Viterbi School of Engineering

Quantum Computing Prognosis

ISI Information Sciences Institute

Quantum computing is a compelling vision

- Has been for a decade now …
- Kinda like fusion energy

Good news

- Potential application space is growing
- Not just Schor's algorithm

In vino veritas

- Analog computing
- Worse …

USC

USC Viterbi School of Engineering

Architecture

ISI Information Sciences Institute

TRIPS (Berger, Texas)

Continuum (Sterling, Louisiana)

Custom vs. Commodity (Dally, Stanford)

Made me wonder …

- **It costs around $20M to design and build custom H/W**
- **Can there be an HPC niche analogous to GPUs?**

USC

Software

Replacing MPI (Gropp, ANL)

- It ain’t gonna be easy ☹

USC Viterbi School of Engineering

Report on S/W and Apps

ISI Information Sciences Institute

- **Commodity Systems**
  - Usual concerns: memory, concurrency, tools, perf., F/T, etc.
- **Specialty architectures (near-term, hybrid?)**
  - Algorithms for unique features
  - Hardware/software co-design
  - Migration path
- **Discontinuities (new architectures such as quantum)**
  - Co-processor model
  - New, custom languages

USC

USC Viterbi School of Engineering

Report on Future H/W

ISI Information Sciences Institute

Near-term future … 10-15 years.

***End-user experience* will drive commercial IT industry**

Beowulfs will exist for science and engineering

- **Can exploit specialized COTS devices (FPGAs, GPUs, etc.)**

Hope for HPCS

Architecture obstructs better algorithms, software, apps.

- **Is it time to revisit COTS-based system assumption?**
- **Barrier to entry for custom H/W is low (~$20M for SOA chip)**

Architecture research pipeline dry (NRC & DSB reports)

- **No new paradigms in ten years**
- **Failure to invest in spite of many reports advocating this**

USC

USC Viterbi School of Engineering

Report on Extreme H/W

ISI Information Sciences Institute

- **The end is near … for Moore's Law at least**
  - **Only reversibly and quantum offer hope beyond kT limit**
  - **Might run out of money first ☹**
- **The future will be heterogeneous**
  - **Emerging technologies might leverage CMOS fab tools**
  - **Fill niches (e.g., memory)**
- **Need roadmap for alternative technologies**
  - **Complement ITRS and proposed NRC architecture maps**
- **Need design challenge competition**
  - **Specify challenge problems**
  - **Compare solutions in different technologies**

USC

USC Viterbi School of Engineering

My Summary

ISI Information Sciences Institute

We *have* need for ZettaFlop/s systems

Si roadmap looks good for another decade

- **Bar is high for alternatives**

There's equivalent inertia in software

- **Rich cluster programming model (e.g., MPI)**

Credible alternatives do exist

- **They'll need varying amounts of time and money to mature**

New technologies may emerge in niches

- **HP's nanoscale memory?**

Specialization in architecture may return

- **Design and H/W costs not unreasonable**

USC